

WHAT IS CLAIMED IS:

1. A method of processing a query command in a distributed computing system in which a plurality of database tables are stored on a plurality of nodes, different portions of at least one database table being stored on at least two of the nodes, the method comprising:

storing a first portion of a first database table and a first portion of a second database table on a first node, and storing a second portion of a first database table and a second portion of a second database table on a second node;

determining a join table definition in response to a query command, said join table definition identifying a subset of said first database table to include in executing said database query command;

generating a first join table from said first portion of said first database table in accordance with said join table definition, and generating a second join table from said second portion of said first database table in accordance with said join table definition;

transmitting said first join table to said second node, and transmitting said second join table to said first node;

comparing said first portion of said second database table with said first join table, and comparing said second portion of said second database table with said second join table to generate a first intermediate results file;

comparing said first portion of said second database table with said second join table, and comparing said second portion of said second database table with said first join table to generate a second intermediate results file; and

generating a final results file from said first intermediate results file and said second intermediate results file.

2. The method of Claim 1, wherein said storing of said first portion of said first database table and said first portion of said second database table on said first node is stored in substantially equal portions.

3. The method of Claim 2, wherein said storing of said first portion of said first database table and said first portion of said second database table on said first node is stored in substantially equal portions according to a round robin distribution.

4. The method of Claim 3, wherein said storing of said second portion of said first database table and said second portion of said second database table on said second node is stored in substantially equal portions according to a round robin distribution.
5. The method of Claim 1, further comprising executing post-processing operations on said final results file.
6. The method of Claim 1, wherein said storing of said first portion of said first database table and said first portion of said second database table is stored on a volatile memory of said first node, and said storing of said second portion of said first database table and said first portion of said second database table is stored on a volatile memory of said second node.
7. The method of Claim 1, further comprising storing said first and second database tables on a persistent storage device.
8. A method of processing a database query command in a distributed database system in which a plurality of database tables are stored on a node having a plurality of logical processors, the method comprising:
 - receiving a database query command;
 - comparing in response to said database query command said first portion of said first database table with said first portion of said second database table, and comparing said first portion of said first database table with said second portion of said second database table to generate a first portion of a results file;
 - comparing in response to said database query command said second portion of said first database table with said first portion of said second database table, and comparing said second portion of said first database table with said second portion of said second database table to generate a second portion of said results file; and
 - executing post-processing operations on said results file to remove duplicate matching records.
9. The method of Claim 8, wherein receiving said database query command comprises receiving a standard query language (SQL) database query command.
10. A distributed database system for processing a database query command in which a plurality of database tables are stored on a plurality of nodes, different portions of at least one database table being stored on at least two of the nodes, the system comprising:

means for storing a first portion of a first database table and a first portion of a second database table on a first node, and storing a second portion of a first database table and a second portion of a second database table on a second node;

means for determining a join table definition in response to a database query command, said join table definition identifying a subset of said first database table to include in executing said database query command;

means for generating a first join table from said first portion of said first database table in accordance with said join table definition, and generating a second join table from said second portion of said first database table in accordance with said join table definition;

means for transmitting said first join table to said second node, and transmitting said second join table to said first node;

means for comparing said first portion of said second database table with said first join table, and comparing said second portion of said second database table with said second join table to generate a first intermediate results file;

means for comparing said first portion of said second database table with said second join table, and comparing said second portion of said second database table with said first join table to generate a second intermediate results file; and

means for generating a final results file from said first intermediate results file and said second intermediate results file.

11. The system of Claim 10, wherein storing of said first portion of said first database table and said first portion of said second database table on said first node comprises storing in substantially equal portions.

12. The system of Claim 11, wherein storing of said first portion of said first database table and said first portion of said second database table on said first node comprises storing in substantially equal portions according to a round robin distribution.

13. The system of Claim 12, wherein storing of said second portion of said first database table and said second portion of said second database table on said second node comprises storing in substantially equal portions according to a round robin distribution.

14. The system of Claim 10, further comprising means for executing post-processing operations on said final results file.

15. The system of Claim 10, wherein said means for storing of said first portion of said first database table and said first portion of said second database table comprises a volatile memory of said first node, and said means for storing of said second portion of said first database table and said first portion of said second database table comprises a volatile memory of said second node.

16. The system of Claim 10, further comprising means for storing said first and second database tables on a persistent storage device.

17. A distributed database system for processing a database query command in which a plurality of database tables are stored on a node having a plurality of logical processors, the method comprising:

means for receiving a database query command;

means for comparing in response to said database query command said first portion of said first database table with said first portion of said second database table, and comparing said first portion of said first database table with said second portion of said second database table to generate a first portion of a results file;

means for comparing in response to said database query command said second portion of said first database table with said first portion of said second database table, and comparing said second portion of said first database table with said second portion of said second database table to generate a second portion of said results file;

and

means for executing post-processing operations on said results file to remove duplicate matching records.

18. The system of Claim 17, wherein means for receiving said database query command comprises means for receiving a standard query language (SQL) database query command.

19. A method of processing a query command in a distributed computing system in which a plurality of database tables are stored on a plurality of nodes, the method comprising:

storing a first database table and a second database table on a first node;

storing a third database table and a fourth database table on a second node;

determining a first join table definition in response to a query command, said first join table definition identifying a subset of said first database table, and generating a second join table definition in response to said query command, said second join table definition identifying a subset of said third database table;

generating a first join table from said first database table in accordance with said first join table definition, and generating a second join table from said third database table in accordance with said second join table definition;

comparing said second database table with said first join table to generate a first intermediate results file;

comparing said fourth database table with said second join table to generate a second intermediate results file; and

generating a final results file from said first intermediate results file and said second intermediate results file.

20. The method of Claim 19, further comprising executing post-processing operations on said final results file.

21. The method of Claim 20, wherein said post-processing operations comprise removing duplicate matching records from said final results file.

22. The method of Claim 19, wherein said first database table and said second database table on said first node are stored on a volatile memory of said first node.

23. The method of Claim 22, wherein said third database table and said fourth database table are stored on a volatile memory of said second node.

24. The method of Claim 19, further comprising storing said first and second database tables on a persistent storage device.

25. The method of Claim 24, further comprising storing said third and fourth database tables on a persistent storage device.